

Intersection Types and Counting*

Paweł Parys

University of Warsaw, Poland

parys@mimuw.edu.pl

We present a new approach to the following meta-problem: given a quantitative property of trees, design a type system such that the desired property for the tree generated by an infinitary ground λ -term corresponds to some property of a derivation of a type for this λ -term, in this type system.

Our approach is presented in the particular case of the language finiteness problem for nondeterministic higher-order recursion schemes (HORSes): given a nondeterministic HORS, decide whether the set of all finite trees generated by this HORS is finite. We give a type system such that the HORS can generate a tree of an arbitrarily large finite size if and only if in the type system we can obtain derivations that are arbitrarily large, in an appropriate sense; the latter condition can be easily decided.

1 Introduction

In this paper we consider λY -calculus, which is an extension of the simply typed λ -calculus by a fixed-point operator Y . A term P of λY -calculus that is of sort¹ o can be used to generate an infinite tree $BT(P)$, called the Böhm tree of P . Trees generated by terms of λY -calculus can be used to faithfully represent the control flow of programs in languages with higher-order functions. Traditionally, Higher Order Recursive Schemes (HORSes) are used for this purpose [8, 12, 17, 16]; this formalism is equivalent to λY -calculus, and the translation between them is rather straightforward [22]. Collapsible Pushdown Systems [10] and Ordered Tree-Pushdown Systems [7] are other equivalent formalisms.

Intersection type systems were intensively used in the context of HORSes, for several purposes like model-checking [13, 16, 5, 21], pumping [14], transformations of HORSes [15, 6], etc. Interestingly, constructions very similar to intersection types were used also on the side of collapsible pushdown systems; they were alternating stack automata [4], and types of stacks [19, 11].

In this paper we show how intersection types can be used for deciding quantitative properties of trees generated by λY -terms. We concentrate on the language finiteness problem for nondeterministic HORSes: given a nondeterministic HORS, decide whether the set of all finite trees generated by this HORS is finite.

This problem can be restated in the world of λY -terms (or standard, deterministic HORSes), generating a single infinite tree. Here, instead of resolving nondeterministic choices during the generation process, we leave them in the resulting tree. Those nondeterministic choices are denoted by a distinguished br (“branch”) symbol, below which we put options that could be chosen. Then to obtain a finite tree generated by the original HORS we just need to recursively choose in every br -labeled node which of the two subtrees we want to consider. Thus, in this setting, the language finiteness problem asks whether the set of all finite trees obtained this way is finite.

*Work supported by the National Science Center (decision DEC-2012/07/D/ST6/02443). A full version of this paper is available [20]

¹We use the word “sort” instead of the usual “type” to avoid confusion with intersection types introduced in this paper.

The difficulty of this problem lies in the fact that sometimes the same finite tree may be found in infinitely many different places of $BT(P)$ (i.e., generated by a nondeterministic HORS in many ways); thus the actual property to decide is whether there is a common bound on the size of each of these trees. This makes the problem inaccessible for standard methods used for analyzing HORSes, as they usually concern only regular properties of the Böhm tree, while boundedness is a problem of different kind. The same difficulty was observed in [14], where they prove a pumping lemma for deterministic HORSes, while admitting (Remark 2.2) that their method is too weak to reason about nondeterministic HORSes.

In order to solve the language finiteness problem, we present an appropriate intersection type system, where derivations are annotated by flags and markers of multiple kinds. The key property of this type system is that the number of flags in a type derivation for a λY -term P approximates the size of some finite tree obtained by resolving nondeterministic choices in the infinite tree $BT(P)$. In consequence, there are type derivations using arbitrarily many flags if, and only if, the answer to the language finiteness problem is “no”.

The language finiteness problem was first attacked in [1] (for safe HORSes only), but their algorithm turned out to be incorrect [2]. To our knowledge, the only known solution of this problem follows from a recent decidability result for the diagonal problem [9, 6]. This problem asks, given a nondeterministic HORS and a set of letters Σ , whether for every $n \in \mathbb{N}$ the HORS generates a finite tree in which every letter from Σ appears at least n times. Clearly, a nondeterministic HORS generates arbitrarily large trees exactly when for some letter a it generates trees having arbitrarily many a letters, i.e., when the answer to the diagonal problem for $\Sigma = \{a\}$ is “yes”.

Our type system is, to some extent, motivated by the algorithm of [6] solving the diagonal problem. This algorithm works by repeating two kinds of transformations of HORSes. The first of them turns the HORS into a HORS generating trees having only a fixed number of branches, one per each letter from Σ (i.e., one branch in our case of $|\Sigma| = 1$). The branches are chosen nondeterministically out of some tree generated by the original HORS; for every $a \in \Sigma$ there is a choice witnessing that a appeared many times in the original tree. Then such a HORS of the special form is turned into a HORS that is of order lower by one, and generates trees having the same nodes as trees generated by the original HORS, but arranged differently (in particular, the new trees may have again arbitrarily many branches). After finitely many repetitions of this procedure, a HORS of order 0 is obtained, and the diagonal problem becomes easily decidable. In some sense we want to do the same, but instead of applying all these transformations one by one, we simulate all of them simultaneously in a single type derivation. In this derivation, for each order n , we allow to place arbitrarily one marker “of order n ”; this corresponds to the nondeterministic choice of one branch in the n -th step of the previous algorithm. We also place some flags “of order n ”, in places that correspond to nodes remaining after the n -th step of the previous algorithm.

The idea of using intersection types for counting is not completely new. Paper [18] presents a type system that, essentially, allows to estimate the size of the β -normal form of a λ -term just by looking at (the number of some flags in) a derivation of a type for this term. A similar idea, but for higher-order pushdown automata, is present in [19], where we can estimate the number of \sharp symbols appearing on a particular, deterministically chosen branch of the generated tree. This previous approach also uses intersection types, where the derivations are marked with just one kind of flags, denoting “productive” places of a λ -term (oppositely to our approach, where we have different flags for different orders, and we also have markers). The trouble with the “one-flag” approach is that it works well only in a completely deterministic setting, where looking independently at each node of the Böhm tree we know how it contributes to the result; the method stops working (or at least we do not know how to prove that it works) in our situation, where we first nondeterministically perform some guesses in the Böhm tree, and only after that we want to count something that depends on the chosen values.

Acknowledgements. I would like to thank Szymon Toruńczyk for stimulating discussions, and anonymous reviewers for useful comments.

2 Preliminaries

Trees. Let Σ be a *ranked alphabet*, i.e., a set of symbols together with a rank function assigning a nonnegative integer to each of the symbols. We assume that Σ contains a distinguished symbol br of rank 2, used to denote nondeterministic choices. A Σ -*labeled tree* is a tree that is rooted (there is a distinguished root node), node-labeled (every node has a label from Σ), ranked (a node with label of rank n has exactly n children), and ordered (children of a node of rank n are numbered from 1 to n).

When t is a Σ -labeled tree t , by $\mathcal{L}(t)$ we denote the set of all finite trees that can be obtained by choosing in every br -labeled node of t which of the two subtrees we want to consider. More formally, we consider the following relation \rightarrow_{br} : we have $t \rightarrow_{\text{br}} u$ if u can be obtained from t by choosing in t a br -labeled node x and its child y , and replacing the subtree starting in x by the subtree starting in y (which removes x and the other subtree of x). Let $\rightarrow_{\text{br}}^*$ be the reflexive transitive closure of \rightarrow_{br} . Then $\mathcal{L}(t)$ contains all trees u that do not use the br label, are finite, and such that $t \rightarrow_{\text{br}}^* u$.

Infinitary λ -calculus. The set of *sorts* (a.k.a. simple types), constructed from a unique basic sort o using a binary operation \rightarrow , is defined as usual. The order of a sort is defined by: $\text{ord}(o) = 0$, and $\text{ord}(\alpha \rightarrow \beta) = \max(1 + \text{ord}(\alpha), \text{ord}(\beta))$.

We consider infinitary, sorted λ -calculus. *Infinitary λ -terms* (or just λ -terms) are defined by coinduction, according to the following rules:

- if $a \in \Sigma$ is a symbol of rank r , and P_1^o, \dots, P_r^o are λ -terms, then $(aP_1^o \dots P_r^o)^o$ is a λ -term,
- for every sort α there are infinitely many variables $x^\alpha, y^\alpha, z^\alpha, \dots$; each of them is a λ -term,
- if $P^{\alpha \rightarrow \beta}$ and Q^α are λ -terms, then $(P^{\alpha \rightarrow \beta} Q^\alpha)^\beta$ is a λ -term, and
- if P^β is a λ -term and x^α is a variable, then $(\lambda x^\alpha. P^\beta)^{\alpha \rightarrow \beta}$ is a λ -term.

We naturally identify λ -terms differing only in names of bound variables. We often omit the sort annotations of λ -terms, but we keep in mind that every λ -term (and every variable) has a particular sort. A λ -term P is *closed* if it has no free variables. Notice that, for technical convenience, a symbol of positive rank is not a λ -term itself, but always comes with arguments. This is not a restriction, since e.g. instead of a unary symbol a one may use the term $\lambda x. ax$.

The order of a λ -term is just the order of its sort. The *complexity* of a λ -term P is the smallest number m such that the order of every subterm of P is at most m . We restrict ourselves to λ -terms that have finite complexity.

A β -reduction is defined as usual. We say that a β -reduction $P \rightarrow_\beta Q$ is of *order* n if it concerns a redex $(\lambda x. R)S$ such that $\text{ord}(\lambda x. R) = n$. In this situation the order of x is at most $n - 1$, but may be smaller (when other arguments of R are of order $n - 1$).

Böhm Trees. We consider Böhm trees only for closed λ -terms of sort o . For such a term P , its *Böhm tree* $BT(P)$ is constructed by coinduction, as follows: if there is a sequence of β -reductions from P to a λ -term of the form $aP_1 \dots P_r$ (where a is a symbol), then the root of the tree t has label a and r children, and the subtree starting in the i -th child is $BT(P_i)$. If there is no sequence of β -reductions from P to a

λ -term of the above form, then $BT(P)$ is the full binary tree with all nodes labeled by br .² By $\mathcal{L}(P)$ we denote $\mathcal{L}(BT(P))$.

λY -calculus. The syntax of λY -calculus is the same as that of finite λ -calculus, extended by symbols $Y^{(\alpha \rightarrow \alpha) \rightarrow \alpha}$, for each sort α . A term of λY -calculus is seen as a term of infinitary λ -calculus if we replace each symbol $Y^{(\alpha \rightarrow \alpha) \rightarrow \alpha}$ by the unique infinite λ -term Z such that Z is syntactically the same as $\lambda x^{\alpha \rightarrow \alpha}.x(Zx)$. In this way, we view λY -calculus as a fragment of infinitary λ -calculus.

It is standard to convert a nondeterministic HORS \mathcal{G} into a closed λY -term P^o such that $\mathcal{L}(P)$ is exactly the set of all finite trees generated by \mathcal{G} . The following theorem, which is our main result, states that the *language finiteness problem* is decidable.

Theorem 1. *Given a closed λY -term P of sort o , one can decide whether $\mathcal{L}(P)$ is finite.*

3 Intersection Type System

In this section we introduce a type system that allows to determine the desired property: whether in $\mathcal{L}(P)$ there is an arbitrarily large tree.

Intuitions. The main novelty of our type system is in using flags and markers, which may label nodes of derivation trees. To every flag and marker we assign a number, called an order. While deriving a type for a λ -term of complexity m , we may place in every derivation tree at most one marker of each order $n \in \{0, \dots, m-1\}$, and arbitrarily many flags of each order $n \in \{0, \dots, m\}$.

Consider first a λ -term M_0 of complexity 0. Such a term actually equals its Böhm tree. Our aim is to describe some finite tree t in $\mathcal{L}(M_0)$, i.e., obtained from M_0 by resolving nondeterministic choices in some way. We thus just put flags of order 0 in all those (appearances of) symbols in M_0 that contribute to this tree t ; the type system ensures that indeed all symbols of some finite tree in $\mathcal{L}(M_0)$ are labeled by a flag. Then clearly we have the desired property that there is a derivation with arbitrarily many flags if, and only if, there are arbitrarily large trees in $\mathcal{L}(M_0)$.

Next, consider a λ -term M_1 that is of complexity 1, and reduces to M_0 . Of course every finite tree from $\mathcal{L}(M_0)$ is composed of symbols appearing already in M_1 ; we can thus already in M_1 label (by order-0 flags) all symbols that contribute to some tree $t \in \mathcal{L}(M_0)$ (and an intersection type system can easily check correctness of such labeling). There is, however, one problem: a single appearance of a symbol in M_1 may result in many appearances in M_0 (since a function may use its argument many times). Due to this, the number of order-0 flags in M_1 does not correspond to the size of t . We rescue ourselves in the following way. In t we choose one leaf, we label it by an order-0 marker, and on the path leading from the root to this marker we place order-1 flags. On the one hand, $\mathcal{L}(M_0)$ contains arbitrarily large trees if, and only if, it contains trees with arbitrarily long paths, i.e., trees with arbitrarily many order-1 flags. On the other hand, we can perform the whole labeling (and the type system can check its correctness) already in M_1 , and the number of order-1 flags in M_1 will be precisely the same as it would be in M_0 . Indeed, in M_1 we have only order-1 functions, i.e., functions that take trees and use them as subtrees of larger trees; although a tree coming as an argument may be duplicated, the order-0 marker can be placed in at most one copy. This means that, while reducing M_1 to M_0 , every symbol of M_1 can result in at most

²Usually one uses a special label \perp of rank 0 for this purpose, but from the perspective of our problem both definitions are equivalent.

one symbol of M_0 lying on the selected path to the order-0 marker (beside of arbitrarily many symbols outside of this path).

This procedure can be repeated for M_2 of complexity 2 that reduces to M_1 via β -reductions of order 2 (and so on for higher orders). We now place a marker of order 1 in some leaf of M_1 ; afterwards, we place an order-2 flag in every node that is on the path to the marked leaf, and that has a child outside of this path whose some descendant is labeled by an order-1 flag. In effect, for some choice of a leaf to be marked, the number of order-2 flags approximates the number of order-1 flags, up to logarithm. Moreover, the whole labeling can be done in M_2 instead of in M_1 , without changing the number of order-2 flags.

In this intuitive description we have talked about labeling “nodes of a λ -term”, but formally we label nodes of a derivation tree that derives a type for the term, in our type system. Every such node contains a type judgment for some subterm of the term.

Type Judgments. For every sort α we define the set \mathcal{T}^α of *types* of sort α , and the set \mathcal{F}^α of *full types* of sort α . This is done as follows, where \mathcal{P} denotes the powerset:

$$\begin{aligned} \mathcal{T}^{\alpha \rightarrow \beta} &= \mathcal{P}(\mathcal{F}_{\text{ord}(\alpha \rightarrow \beta)}^\alpha) \times \mathcal{T}^\beta, & \mathcal{T}^o &= o, \\ \mathcal{F}_k^\alpha &= \{(k, F, M, \tau) \mid F, M \subseteq \{0, \dots, k-1\}, F \cap M = \emptyset, \tau \in \mathcal{T}^\alpha\}, & \mathcal{F}^\alpha &= \bigcup_{k \in \mathbb{N}} \mathcal{F}_k^\alpha. \end{aligned}$$

Notice that the sets \mathcal{T}^α and \mathcal{F}_k^α are finite (unlike \mathcal{F}^α). A type $(T, \tau) \in \mathcal{T}^{\alpha \rightarrow \beta}$ is denoted as $T \rightarrow \tau$. A full type $\hat{\tau} = (k, F, M, \tau) \in \mathcal{F}_k^\alpha$ consists of its order k , a set F of flag orders, a set M of marker orders, and a type τ ; we write $\text{ord}(\hat{\tau}) = k$. In order to distinguish types from full types, the latter are denoted by letters with a hat, like $\hat{\tau}$.

A *type judgment* is of the form $\Gamma \vdash P : \hat{\tau} \triangleright c$, where Γ , called a *type environment*, is a function that maps every variable x^α to a subset of \mathcal{F}^α , P is a λ -term, $\hat{\tau}$ is a full type of the same sort as P (i.e., $\hat{\tau} \in \mathcal{F}^\beta$ when P is of sort β), and $c \in \mathbb{N}$.

As usual for intersection types, the intuitive meaning of a type $T \rightarrow \tau$ is that a λ -term having this type can return a λ -term having type τ , while taking an argument for which we can derive all full types from T . Moreover, in \mathcal{T}^o there is just one type o , which can be assigned to every λ -term of sort o . Suppose that we have derived a type judgment $\Gamma \vdash P : \hat{\tau} \triangleright c$ with $\hat{\tau} = (m, F, M, \tau)$. Then

- τ is the type derived for P ;
- Γ contains full types that could be used for free variables of P in the derivation;
- m bounds the order of flags and markers that could be used in the derivation: flags could be of order at most m , and markers of order at most $m-1$;
- $M \subseteq \{0, \dots, m-1\}$ contains the orders of markers used in the derivation, together with those provided by free variables (i.e., we imagine that some derivations, specified by the type environment, are already substituted in our derivation for free variables); we, however, do not include markers provided by arguments of the term (i.e., coming from the sets T_i when $\tau = T_1 \rightarrow \dots \rightarrow T_k \rightarrow o$);
- F contains those numbers $n \in \{0, \dots, m-1\}$ (excluding $n = m$) for which a flag of order n is placed in the derivation itself, or provided by a free variable, or provided by an argument; for technical convenience we, however, remove n from F whenever $n \in M$ (when $n \in M$, the information about order- n flags results in placing an order- $(n+1)$ flag, and need not to be further propagated);
- c , called a *flag counter*, counts the number of order- m flags present in the derivation.

Type System. Before giving rules of the type system, we need a few definitions. We use the symbol \uplus to denote disjoint union. When $A \subseteq \mathbb{N}$ and $n \in \mathbb{N}$, we write $A \upharpoonright_{<n}$ for $\{k \in A \mid k < n\}$, and similarly $A \upharpoonright_{\geq n}$ for $\{k \in A \mid k \geq n\}$. By ε we denote the type environment mapping every variable to \emptyset , and by $\Gamma[x \mapsto T]$ the type environment mapping x to T and every other variable y to $\Gamma(y)$.

Let us now say how a type environment Γ from the conclusion of a rule may be split into type environments $(\Gamma_i)_{i \in I}$ used in premisses of the rule: we say that $\text{Split}(\Gamma \mid (\Gamma_i)_{i \in I})$ holds if and only if for every variable x it holds $\Gamma_i(x) \subseteq \Gamma(x)$ for every $i \in I$, and every full type from $\Gamma(x)$ providing some markers (i.e., (k, F, M, τ) with $M \neq \emptyset$) appears in some $\Gamma_i(x)$. Full types with empty M may be discarded and duplicated freely. This definition forbids to discard full types with nonempty M , and from elsewhere it will follow that they cannot be duplicated. As a special case $\text{Split}(\Gamma \mid \Gamma')$ describes how a type environment can be weakened.

All type derivations are assumed to be finite (although we derive types mostly for infinite λ -terms, each type derivation analyzes only a finite part of a term). Rules of the type system will guarantee that the order m of derived full types will be the same in the whole derivation (although in type environments there may be full types of different orders).

We are ready to give the first three rules of our type system:

$$\frac{\Gamma \vdash P_i : \hat{\tau} \triangleright c \quad i \in \{1, 2\}}{\Gamma \vdash \text{br } P_1 P_2 : \hat{\tau} \triangleright c} \text{ (BR)} \quad \frac{\text{Split}(\Gamma \mid \varepsilon[x \mapsto \{(k, F, M', \tau)\}]) \quad M \upharpoonright_{<k} = M'}{\Gamma \vdash x : (m, F, M, \tau) \triangleright 0} \text{ (VAR)}$$

$$\frac{\Gamma'[x \mapsto T] \vdash P : (m, F, M, \tau) \triangleright c \quad \text{Split}(\Gamma \mid \Gamma') \quad \Gamma'(x) = \emptyset}{\Gamma \vdash \lambda x. P : (m, F, M \setminus \bigcup_{(k, F', M', \sigma) \in T} M', T \rightarrow \tau) \triangleright c} (\lambda)$$

We see that to derive a type for the nondeterministic choice $\text{br } P_1 P_2$, we need to derive it either for P_1 or for P_2 .

The (VAR) rule allows to have in the resulting set M some numbers that do not come from the set M' assigned to x by the type environment; these are the orders of markers placed in the leaf using this rule. Notice, however, that we allow here only orders not smaller than k (which is the order of the superterm $\lambda x. P$ binding this variable x). This is consistent with the intuitive description of the type system (page 51), which says that a marker of order n can be put in a place that will be a leaf after performing all β -reductions of orders greater than n . Indeed, the variable x remains a leaf after performing β -reductions of orders greater than k , but while performing β -reductions of order k this leaf will be replaced by a subterm substituted for x . Recall also that, by definition of a type judgment, we require that $(k, F, M', \tau) \in \mathcal{F}_k^\alpha$ and $(m, F, M, \tau) \in \mathcal{F}_m^\alpha$, for appropriate sort α ; this introduces a bound on maximal numbers that may appear in the sets F and M .

Example 1. Denoting $\hat{\rho}_1 = (1, \emptyset, \{0\}, o)$ we can derive:

$$\frac{}{\varepsilon[x \mapsto \{\hat{\rho}_1\}] \vdash x : (2, \emptyset, \{0\}, o) \triangleright 0} \text{ (VAR)} \quad \frac{}{\varepsilon[x \mapsto \{\hat{\rho}_1\}] \vdash x : (2, \emptyset, \{0, 1\}, o) \triangleright 0} \text{ (VAR)}$$

In the derivation on the right, the marker of order 1 is placed in the conclusion of the rule.

The (λ) rule allows to use (in a subderivation concerning the λ -term P) the variable x with all full types given in the set T . When the sort of $\lambda x. P$ is $\alpha \rightarrow \beta$, by definition of $\mathcal{T}^{\alpha \rightarrow \beta}$ we have that all full types in T have the same order $k = \text{ord}(\alpha \rightarrow \beta)$ (since $(T \rightarrow \tau) \in \mathcal{T}^{\alpha \rightarrow \beta}$). Recall that we intend to store in the set M the markers contained in the derivation itself and those provided by free variables, but not those provided by arguments. Because of this, in the conclusion of the rule we remove from M the markers provided by x . This operation makes sense only because there is at most one marker of each order, so

markers provided by x cannot be provided by any other free variable nor placed in the derivation itself. The set F , unlike M , stores also flags provided by arguments, so we do not need to remove anything from F .

Example 2. The (λ) rule can be used, e.g., in the following way (where a is a symbol of rank 1):

$$\frac{\varepsilon[x \mapsto \{\hat{\rho}_1\}] \vdash ax : (2, \{1\}, \{0\}, o) \triangleright 0}{\varepsilon \vdash \lambda x. ax : (2, \{1\}, \emptyset, \{\hat{\rho}_1\} \rightarrow o) \triangleright 0} (\lambda) \quad \frac{\varepsilon[x \mapsto \{\hat{\rho}_1\}] \vdash ax : (2, \emptyset, \{0, 1\}, o) \triangleright 1}{\varepsilon \vdash \lambda x. ax : (2, \emptyset, \{1\}, \{\hat{\rho}_1\} \rightarrow o) \triangleright 1} (\lambda)$$

Notice that in the conclusion of the rule, in both examples, we remove 0 from the set of marker orders, because the order-0 marker is provided by x .

The next two rules use a predicate $Comp_m$, saying how flags and markers from premisses contribute to the conclusion. It takes “as input” pairs (F_i, c_i) for $i \in I$; each of them consists of the set of flag orders F_i and of the flag counter c_i from some premiss. Moreover, the predicate takes a set of marker orders M from the current type judgment (it contains orders of markers used in the derivation, including those provided by free variables). The goal is to compute the set of flag orders F and the flag counter c that should be placed in the current type judgment. First, for each $n \in \{1, \dots, m\}$ consecutively, we decide whether a flag of order n should be placed on the current type judgment. We follow here the rules mentioned in the intuitive description. Namely, we place a flag of order n if we are on the path leading to the marker of order $n-1$ (i.e., if $n-1 \in M$), and simultaneously we receive an information about a flag of order $n-1$. By receiving this information we mean that either a flag of order $n-1$ was placed on the current type judgment, or $n-1$ belongs to some set F_i . Actually, we place multiple flags of order n : one per each flag of order $n-1$ placed on the current type judgment, and one per each set F_i containing $n-1$. Then, we compute F and c . In c we store the number of flags of the maximal order m : we sum all the numbers c_i , and we add the number of order- m flags placed on the current type judgment. In F we keep elements of all F_i , and we add the orders n of flags that were placed on the current type judgment. We, however, remove from F all elements of M . This is because every flag of some order $n-1$ should result in creating at most one flag of order n , in the closest ancestor that lies on the path leading to the marker of order $n-1$. If we have created an order- n flag on the current type judgment, i.e., if $n-1 \in M$, we do not want to do this again in the parent.

Below we give a formal definition, in which f'_n contains the number of order- n flags placed on the current type judgment, while f_n additionally counts the number of premisses for which $n \in F_i$. We say that $Comp_m(M; ((F_i, c_i))_{i \in I}) = (F, c)$ when

$$F = \{n \in \{0, \dots, m-1\} \mid f_n > 0 \wedge n \notin M\}, \quad c = f'_m + \sum_{i \in I} c_i, \quad \text{where, for } n \in \{0, \dots, m\},$$

$$f_n = f'_n + \sum_{i \in I} |F_i \cap \{n\}|, \quad f'_n = \begin{cases} f'_{n-1} & \text{if } n-1 \in M, \\ 0 & \text{otherwise.} \end{cases}$$

We now present a rule for constants other than br:

$$\frac{\begin{array}{l} \Gamma_i \vdash P_i : (m, F_i, M_i, o) \triangleright c_i \text{ for each } i \in \{1, \dots, r\} \quad M = M' \uplus M_1 \uplus \dots \uplus M_r \\ (m=0) \Rightarrow (F' = \emptyset \wedge c' = 1) \quad (m>0) \Rightarrow (F' = \{0\} \wedge c' = 0) \quad (r>0) \Rightarrow (M' = \emptyset) \\ a \neq \text{br} \quad \text{Split}(\Gamma \mid \Gamma_1, \dots, \Gamma_r) \quad \text{Comp}_m(M; (F', c'), (F_1, c_1), \dots, (F_r, c_r)) = (F, c) \end{array}}{\Gamma \vdash aP_1 \dots P_r : (m, F, M, o) \triangleright c} (\text{CON})$$

Here, the conditions in the second line say that in a node using the (CON) rule we always place a flag of order 0 (via F' or via c' , depending on m), and that if the node is a leaf (i.e., $r = 0$), then we

are allowed to place markers of arbitrary order (via M'). Then to the $Comp_m$ predicate, beside of pairs (F_i, c_i) coming from premisses, we also pass the information (F', c') about the order-0 flag placed in the current node; this predicate decides whether we should place also some flags of positive orders. Let us emphasize that in this rule (and similarly in the next rule) we have a disjoint union $M' \uplus M_1 \uplus \dots \uplus M_r$, which ensures that a marker of any order may be placed only in one node of a derivation.

Example 3. The (CON) rule may be instantiated in the following way:

$$\frac{\varepsilon[x \mapsto \{\hat{\rho}_1\}] \vdash x : (2, \emptyset, \{0\}, o) \triangleright 0}{\varepsilon[x \mapsto \{\hat{\rho}_1\}] \vdash ax : (2, \{1\}, \{0\}, o) \triangleright 0} \text{ (CON)} \quad \frac{\varepsilon[x \mapsto \{\hat{\rho}_1\}] \vdash x : (2, \emptyset, \{0, 1\}, o) \triangleright 0}{\varepsilon[x \mapsto \{\hat{\rho}_1\}] \vdash ax : (2, \emptyset, \{0, 1\}, o) \triangleright 1} \text{ (CON)}$$

In the left example, flags of order 0 and 1 are placed in the conclusion of the rule (a flag of order 0 is created because we are in a constant; since the marker of order 0 is visible, we do not put 0 into the set of flag orders, but instead we create a flag of order 1). In the right example, a marker of order 1 is visible, which causes that this time flags of order 0, 1, and 2 are placed in the conclusion of the (CON) rule (again, we do not put 0 nor 1 into the set of flag orders, because of 0 and 1 in the set of marker orders).

The next rule describes application:

$$\frac{\begin{array}{c} \Gamma' \vdash P : (m, F', M', \{(ord(P), F_i \upharpoonright_{<ord(P)}, M_i \upharpoonright_{<ord(P)}, \tau_i) \mid i \in I\} \rightarrow \tau) \triangleright c' \\ \Gamma_i \vdash Q : (m, F_i, M_i, \tau_i) \triangleright c_i \text{ for each } i \in I \quad M = M' \uplus \biguplus_{i \in I} M_i \\ ord(P) \leq m \quad Split(\Gamma \mid \Gamma', (\Gamma_i)_{i \in I}) \quad Comp_m(M; (F', c'), ((F_i \upharpoonright_{\geq ord(P)}, c_i))_{i \in I}) = (F, c) \end{array}}{\Gamma \vdash PQ : (m, F, M, \tau) \triangleright c} \text{ (@)}$$

In this rule, it is allowed (but in fact useless) that for two different $i \in I$ the full types (m, F_i, M_i, τ_i) are equal. It is also allowed that $I = \emptyset$, in which case no type needs to be derived for Q . Observe how flags and markers coming from premisses concerning Q are propagated: only flags and markers of order $n < ord(P)$ are visible to P , while only flags of order $n \geq ord(P)$ are passed to the $Comp_m$ predicate. This can be justified if we recall the intuitions staying behind the type system (see page 51). Indeed, while considering flags and markers of order n , we should imagine the λ -term obtained from the current λ -term by performing all β -reductions of all orders greater than n ; the distribution of flags and markers of order n in the current λ -term actually simulates their distribution in this imaginary λ -term. Thus, if $n < ord(P)$, then our application will disappear in this imaginary λ -term, and Q will be already substituted somewhere in P ; for this reason we need to pass the information about flags and markers of order n from Q to P . Conversely, if $n \geq ord(P)$, then in the imaginary λ -term the considered application will be still present, and in consequence the subterm corresponding to P will not see flags and markers of order n placed in the subterm corresponding to Q .

Example 4. Denote by $\hat{\tau}_f$ and $\hat{\tau}_m$ the types derived in Example 2:

$$\hat{\tau}_f = (2, \{1\}, \emptyset, \{\hat{\rho}_1\} \rightarrow o), \quad \text{and} \quad \hat{\tau}_m = (2, \emptyset, \{1\}, \{\hat{\rho}_1\} \rightarrow o).$$

Then, using the (@) rule, we can derive (where e is a symbol of rank 0, and f a variable):

$$\frac{\frac{\varepsilon[f \mapsto \{\hat{\tau}_m\}] \vdash f : \hat{\tau}_m \triangleright 0 \text{ (VAR)}}{\varepsilon[f \mapsto \{\hat{\tau}_f, \hat{\tau}_m\}] \vdash f e : (2, \emptyset, \{0, 1\}, o) \triangleright 1} \text{ (CON)} \quad \frac{\varepsilon \vdash e : (2, \{1\}, \{0\}, o) \triangleright 0}{\varepsilon[f \mapsto \{\hat{\tau}_f, \hat{\tau}_m\}] \vdash f e : (2, \emptyset, \{0, 1\}, o) \triangleright 1} \text{ (@)}$$

Recall that $\hat{\rho}_1 = (1, \emptyset, \{0\}, o)$. In the conclusion of the (@) rule the information about a flag of order 1 (from the second premiss) meets the information about the marker of order 1 (from the first premiss), and

thus a flag of order 2 is placed, which increases the flag counter. Notice that we have discarded the full type $\hat{\tau}_f$ assigned to f in the type environment; this is allowed because $\hat{\tau}_f$ provides no markers (equally well $\hat{\tau}_f$ could be assigned to f also in one or two of the premisses, and discarded there). On the other hand, the full type $\hat{\tau}_m$ provides markers, so it cannot be discarded nor duplicated (in particular, we could not pass it to the conclusion of the (CON) rule).

The key property of the type system is described by the following theorem.

Theorem 2. *Let P be a closed λ -term of sort o and complexity m . Then $\mathcal{L}(P)$ is infinite if and only if for arbitrarily large c we can derive $\varepsilon \vdash P : \hat{\rho}_m \triangleright c$, where $\hat{\rho}_m = (m, \emptyset, \{0, \dots, m-1\}, o)$.*

The left-to-right implication of Theorem 2 (completeness of the type system) is shown in Section 4, while the opposite implication (soundness of the type system) in Section 5. In Section 6 we discuss how Theorem 1 follows from Theorem 2. Before all that, we give a few more examples of derivations, illustrating the type system and Theorem 2.

Example 5. In this example we analyze the λ -term $P_1 = R(\lambda x.ax)$, where R is defined by coinduction as $R = (\lambda f.br(fe)(R(\lambda x.f(fx))))$. As previously, a and e are symbols of rank 1 and 0, respectively. In $\mathcal{L}(P_1)$ there are trees that consist of a branch of a symbols ended with an e symbol, but only those where the number of a symbols is 2^k for some $k \in \mathbb{N}$. Notice that the complexity of P_1 is 2.

Continuing Example 4, we derive the full type $\hat{\sigma}_R = (2, \emptyset, \{0\}, \{\hat{\tau}_f, \hat{\tau}_m\} \rightarrow o)$ for R :

$$\frac{\frac{\varepsilon[f \mapsto \{\hat{\tau}_f, \hat{\tau}_m\}] \vdash fe : (2, \emptyset, \{0, 1\}, o) \triangleright 1}{\varepsilon[f \mapsto \{\hat{\tau}_f, \hat{\tau}_m\}] \vdash br(fe)(R(\lambda x.f(fx))) : (2, \emptyset, \{0, 1\}, o) \triangleright 1} \text{ (BR)}}{\varepsilon \vdash R : \hat{\sigma}_R \triangleright 1} \text{ (\lambda)}$$

Next, we derive the same full type for R , but using the second argument of the br symbol; this results in greater values of the flag counter. We start by deriving the full type $\hat{\tau}_f$ for the subterm $\lambda x.f(fx)$:

$$\frac{\frac{\varepsilon[f \mapsto \{\hat{\tau}_f\}] \vdash f : \hat{\tau}_f \triangleright 0 \quad \frac{\varepsilon[x \mapsto \{\hat{\rho}_1\}] \vdash x : (2, \emptyset, \{0\}, o) \triangleright 0}{\varepsilon[f \mapsto \{\hat{\tau}_f\}, x \mapsto \{\hat{\rho}_1\}] \vdash fx : (2, \{1\}, \{0\}, o) \triangleright 0} \text{ (@)}}{\varepsilon[f \mapsto \{\hat{\tau}_f\}] \vdash f : \hat{\tau}_f \triangleright 0} \text{ (@)} \\ \frac{\varepsilon[f \mapsto \{\hat{\tau}_f\}, x \mapsto \{\hat{\rho}_1\}] \vdash f(fx) : (2, \{1\}, \{0\}, o) \triangleright 0}{\varepsilon[f \mapsto \{\hat{\tau}_f\}] \vdash \lambda x.f(fx) : \hat{\tau}_f \triangleright 0} \text{ (\lambda)}$$

In the above derivation there are no flags nor markers. Next, we derive $\hat{\tau}_m$ for the same subterm:

$$\frac{\frac{\varepsilon[f \mapsto \{\hat{\tau}_m\}] \vdash f : \hat{\tau}_m \triangleright 0 \quad \frac{\varepsilon[x \mapsto \{\hat{\rho}_1\}] \vdash x : (2, \emptyset, \{0\}, o) \triangleright 0}{\varepsilon[f \mapsto \{\hat{\tau}_m\}, x \mapsto \{\hat{\rho}_1\}] \vdash fx : (2, \emptyset, \{0, 1\}, o) \triangleright 0} \text{ (@)}}{\varepsilon[f \mapsto \{\hat{\tau}_f\}] \vdash f : \hat{\tau}_f \triangleright 0} \text{ (@)} \\ \frac{\varepsilon[f \mapsto \{\hat{\tau}_f, \hat{\tau}_m\}, x \mapsto \{\hat{\rho}_1\}] \vdash f(fx) : (2, \emptyset, \{0, 1\}, o) \triangleright 1}{\varepsilon[f \mapsto \{\hat{\tau}_f, \hat{\tau}_m\}] \vdash \lambda x.f(fx) : \hat{\tau}_m \triangleright 1} \text{ (\lambda)}$$

Below the lower (@) rule the information about a flag of order 1 meets the information about the marker of order 1, and thus a flag of order 2 is placed, which increases the flag counter. We continue with the λ -term R :

$$\frac{\frac{\varepsilon \vdash R : \hat{\sigma}_R \triangleright c \quad \varepsilon[f \mapsto \{\hat{\tau}_f\}] \vdash \lambda x.f(fx) : \hat{\tau}_f \triangleright 0 \quad \varepsilon[f \mapsto \{\hat{\tau}_f, \hat{\tau}_m\}] \vdash \lambda x.f(fx) : \hat{\tau}_m \triangleright 1}{\varepsilon[f \mapsto \{\hat{\tau}_f, \hat{\tau}_m\}] \vdash R(\lambda x.f(fx)) : (2, \emptyset, \{0, 1\}, o) \triangleright c+1} \text{ (@)}}{\frac{\varepsilon[f \mapsto \{\hat{\tau}_f, \hat{\tau}_m\}] \vdash br(fe)(R(\lambda x.f(fx))) : (2, \emptyset, \{0, 1\}, o) \triangleright c+1}{\varepsilon \vdash R : \hat{\sigma}_R \triangleright c+1} \text{ (BR)}} \text{ (\lambda)}$$

In this fragment of a derivation no flag nor marker is placed. In particular, there is no order-2 flag in conclusion of the (@) rule, although its second premiss provides a flag of order 1 while the third premiss provides the marker of order 1. We recall from the definition of the (@) rule that the information about flags and markers coming from the arguments is divided into two parts. Numbers smaller than the order of the operator ($ord(R) = 2$ in our case) are passed to the operator, while only greater numbers (≥ 2 in our case) contribute in creating new flags via the *Comp* predicate.

By composing the above fragments of a derivation, we can derive $\varepsilon \vdash R : \hat{\sigma}_R \triangleright c$ for every $c \geq 1$. Recall that in Examples 1-3 we have derived $\varepsilon \vdash \lambda x.ax : \hat{\tau}_f \triangleright 0$ and $\varepsilon \vdash \lambda x.ax : \hat{\tau}_m \triangleright 1$. Together with the above, this allows to derive for P_1 the full type $\hat{\rho}_2 = (2, \emptyset, \{0, 1\}, o)$ (appearing in Theorem 2):

$$\frac{\varepsilon \vdash R : \hat{\sigma}_R \triangleright c \quad \varepsilon \vdash \lambda x.ax : \hat{\tau}_f \triangleright 0 \quad \varepsilon \vdash \lambda x.ax : \hat{\tau}_m \triangleright 1}{\varepsilon \vdash P_1 : \hat{\rho}_2 \triangleright c + 1} (@)$$

We can notice a correspondence between a derivation with flag counter $c + 1$ and a tree in $\mathcal{L}(P)$ of size $2^{c-1} + 1$. We remark that in every of these derivations only three flags of order 0 and only three flags of order 1 are present, in the three nodes using the (CON) rule.

Example 6. Consider a similar λ -term $P_2 = R(\lambda x.bxx)$, where R is as previously, and b is a symbol of rank 2. In $\mathcal{L}(P_2)$ we have, for every $k \in \mathbb{N}$, a full binary tree in which every branch consist of 2^k symbols b and ends with an e symbol.

This time for the subterm $\lambda x.bxx$ we need to derive three full types:

$$\begin{aligned} \hat{\tau}'_0 &= (2, \{0\}, \emptyset, \{(1, \{0\}, \emptyset, o) \rightarrow o\}), \\ \hat{\tau}'_f &= (2, \{1\}, \emptyset, \{(1, \{0\}, \emptyset, o), \hat{\rho}_1 \} \rightarrow o), \quad \text{and} \\ \hat{\tau}'_m &= (2, \emptyset, \{1\}, \{(1, \{0\}, \emptyset, o), \hat{\rho}_1 \} \rightarrow o). \end{aligned}$$

The last one is derived with flag counter 1. Notice that $\hat{\tau}'_f$ and $\hat{\tau}'_m$ need now two full types for the argument x ; the new one $(1, \{0\}, \emptyset, o)$ describes the subtree that is not on the path to the order-0 marker. We also have a new full type $\hat{\tau}'_0$ that describes the use of $\lambda x.bxx$ outside of the path to the order-0 marker.

Then, similarly as in the previous example, for every $c \geq 1$ we can derive $\varepsilon \vdash R : \hat{\sigma}'_R \triangleright c$, where $\hat{\sigma}'_R = (2, \emptyset, \{0\}, \{\hat{\tau}'_0, \hat{\tau}'_f, \hat{\tau}'_m\} \rightarrow o)$. Again, this allows to derive $\varepsilon \vdash P_2 : \hat{\rho}_2 \triangleright c + 1$. This time a derivation with flag counter $c + 1$ corresponds to a tree in $\mathcal{L}(P)$ of size $2^h - 1$ with $h = 2^{c-1} + 1$.

Example 7. Next, consider the λ -term $P_3 = R(\lambda x.x)$. The only tree in $\mathcal{L}(P_3)$ consists of a single e node. Let us see how the derivation from Example 5 has to be modified. The full type $\hat{\tau}_m$ can still be derived for $\lambda x.x$ (although with flag counter 0 now), but instead of $\hat{\tau}_f$ we have to use $\hat{\tau}''_f = (2, \emptyset, \emptyset, \{\hat{\rho}_1\} \rightarrow o)$ that provides no flag of order 1:

$$\frac{\varepsilon[x \mapsto \{\hat{\rho}_1\}] \vdash x : (2, \emptyset, \{0\}, o) \triangleright 0}{\varepsilon \vdash \lambda x.x : \hat{\tau}''_f \triangleright 0} \begin{matrix} (\text{VAR}) \\ (\lambda) \end{matrix} \quad \frac{\varepsilon[x \mapsto \{\hat{\rho}_1\}] \vdash x : (2, \emptyset, \{0, 1\}, o) \triangleright 0}{\varepsilon \vdash \lambda x.x : \hat{\tau}_m \triangleright 0} \begin{matrix} (\text{VAR}) \\ (\lambda) \end{matrix}$$

Next, for R we want to derive the full type $\hat{\sigma}''_R = (2, \emptyset, \{0\}, \{\hat{\tau}''_f, \hat{\tau}_m\} \rightarrow o)$. We can easily adopt every of the previous derivations for $\varepsilon \vdash R : \hat{\sigma}_R \triangleright c$: we basically replace every $\hat{\tau}_f$ by $\hat{\tau}''_f$. The key point is that while deriving the full type $\hat{\tau}_m$ for the subterm $\lambda x.f(fx)$, previously in the lower (@) rule we have received information about an order-1 flag, and thus we have created an order-2 flag and increased the flag counter; this time there is no information about an order-1 flag, and thus we do not create an order-2 flag and do not increase the flag counter. In consequence, even if this part of the derivation is repeated arbitrarily many times, the value of the flag counter of the whole derivation remains 1.

Example 8. Finally, consider the λ -term $P_4 = (\lambda g.P_3)(\lambda x.a(a(\dots(ax)\dots))$, which β -reduces to P_3 . Notice that we can create the following derivation:

$$\frac{\frac{\frac{\varepsilon[x \mapsto \{\hat{\rho}_1\}] \vdash x : (2, \emptyset, \{0\}, o) \triangleright 0 \quad (\text{VAR})}{\varepsilon[x \mapsto \{\hat{\rho}_1\}] \vdash ax : (2, \{1\}, \{0\}, o) \triangleright 0 \quad (\text{CON})}}{\varepsilon[x \mapsto \{\hat{\rho}_1\}] \vdash a(a(\dots(ax)\dots)) : (2, \{1\}, \{0\}, o) \triangleright 0 \quad (\text{CON})} \quad \vdots \quad \frac{\varepsilon[x \mapsto \{\hat{\rho}_1\}] \vdash a(a(\dots(ax)\dots)) : (2, \{1\}, \{0\}, o) \triangleright 0 \quad (\text{CON})}{\varepsilon \vdash \lambda x.a(a(\dots(ax)\dots)) : \hat{\tau}_f \triangleright 0 \quad (\lambda)}$$

Every (CON) rule used in this derivation places in its conclusion an order-0 flag and an order-1 flag. This derivation can be used as a part of a derivation for P_4 :

$$\frac{\frac{\varepsilon[g \mapsto \{\hat{\tau}_f\}] \vdash P_3 : \hat{\rho}_2 \triangleright 1}{\varepsilon \vdash \lambda g.P_3 : (2, \emptyset, \{0, 1\}, \{\hat{\tau}_f\} \rightarrow o) \triangleright 1 \quad (\lambda)} \quad \varepsilon \vdash \lambda x.a(a(\dots(ax)\dots)) : \hat{\tau}_f \triangleright 0 \quad (@)}{\varepsilon \vdash P_4 : \hat{\rho}_2 \triangleright 1}$$

Because $\hat{\tau}_f$ provides no markers, it can be removed from the type environment and thus for P_3 we can use the derivation from the previous example. We thus obtain a derivation for P_4 in which there are many order-0 and order-1 flags (but only one flag of order 2). This shows that in the flag counter we indeed need to count only the number of flags of the maximal order (not, say, the total number of flags of all orders).

4 Completeness

The proof of the left-to-right implication of Theorem 2 is divided into the following three lemmata. Recall that a β -reduction $P \rightarrow_\beta Q$ is of order n if it concerns a redex $(\lambda x.R)S$ such that $\text{ord}(\lambda x.R) = n$. The number of nodes of a tree t is denoted $|t|$. As in Theorem 2, we denote $\hat{\rho}_m = (m, \emptyset, \{0, \dots, m-1\}, o)$.

Lemma 3. *Let P be a closed λ -term of sort o and complexity m , and let $t \in \mathcal{L}(P)$. Then there exist λ -terms Q_m, Q_{m-1}, \dots, Q_0 such that $P = Q_m$, and for every $k \in \{1, \dots, m\}$ the term Q_{k-1} can be reached from Q_k using only β -reductions of order k , and we can derive $\varepsilon \vdash Q_0 : \hat{\rho}_0 \triangleright |t|$.*

Lemma 4. *Suppose that we can derive $\varepsilon \vdash P : \hat{\rho}_m \triangleright c$. Then we can also derive $\varepsilon \vdash P : \hat{\rho}_{m+1} \triangleright c'$ for some $c' \geq \log_2 c$.*

Lemma 5. *Suppose that $P \rightarrow_\beta Q$ is a β -reduction of order m , and we can derive $\Gamma \vdash Q : \hat{\tau} \triangleright c$ with $\text{ord}(\hat{\tau}) = m$. Then we can also derive $\Gamma \vdash P : \hat{\tau} \triangleright c$.*

Now the left-to-right implication of Theorem 2 easily follows. Indeed, take a closed λ -term P of sort o and complexity m such that $\mathcal{L}(P)$ is infinite, and take any $c \in \mathbb{N}$. By \log_2^k we denote the k -fold application of the logarithm: $\log_2^0 x = x$ and $\log_2^{k+1} x = \log_2(\log_2^k x)$. Since $\mathcal{L}(P)$ is infinite, it contains a tree t so big that $\log_2^m |t| \geq c$. We apply Lemma 3 to this tree, obtaining λ -terms Q_m, Q_{m-1}, \dots, Q_0 and a derivation of $\varepsilon \vdash Q_0 : \hat{\rho}_0 \triangleright |t|$. Then repeatedly for every $k \in \{1, \dots, m\}$ we apply Lemma 4, obtaining a derivation of $\varepsilon \vdash Q_{k-1} : \hat{\rho}_k \triangleright c_k$ for some $c_k \geq \log_2^k |t|$, and Lemma 5 for every β -reduction (of order k) between Q_k and Q_{k-1} , obtaining a derivation of $\varepsilon \vdash Q_k : \hat{\rho}_k \triangleright c_k$. We end with a derivation of $\varepsilon \vdash P : \hat{\rho}_m \triangleright c_m$, where $c_m \geq \log_2^m |t| \geq c$, as needed. In the remaining part of this section we prove the three lemmata.

Proof of Lemma 3 (sketch). Recall that $t \in \mathcal{L}(P)$ is a finite tree, thus it can be found in some finite prefix of the Böhm tree of P . By definition, this prefix will be already expanded after performing some finite number of β -reductions from P . We need to observe that these β -reductions can be rearranged, so that those of higher order are performed first. The key point is to observe that when we perform a β -reduction of some order k , then no new β -redexes of higher order appear in the term. Indeed, suppose that $(\lambda x.R)S$ is changed into $R[S/x]$ somewhere in a term, where $\text{ord}(\lambda x.R) = k$. One new redex that may appear is when R starts with a λ , and to the whole $R[S/x]$ some argument is applied; this redex is of order $\text{ord}(R) \leq k$. Some other redexes may appear when S starts with a λ , and is substituted for such appearance of x to which some argument is applied; but this redex is of order $\text{ord}(S) < k$.

We can thus find a sequence of β -reductions in which β -reductions are arranged according to their order, that leads from P to some Q_0 such that t can be found in the prefix of Q_0 that is already expanded to a tree. It is now a routine to use the rules of our type system and derive $\varepsilon \vdash Q_0 : \hat{\rho}_0 \triangleright |t|$: in every br-labeled node we choose the subtree in which t continues, and this effects in counting the number of nodes of t in the flag counter. \square

Proof of Lemma 4. Consider some derivation of $\varepsilon \vdash P : \hat{\rho}_m \triangleright c$. In this derivation we choose a leaf in which we will put the order- m marker, as follows. Starting from the root of the derivation, we repeatedly go to this premiss in which the flag counter is the greatest (arbitrarily in the case of a tie). In every node that is not on the path to the selected leaf, we replace the current type judgment $\Gamma \vdash Q : (m, F, M, \tau) \triangleright d$ by $\Gamma \vdash Q : (m+1, F', M, \tau) \triangleright 0$, where $F' = F \cup \{m\}$ if $d > 0$, and $F' = F$ otherwise. In the selected leaf and all its ancestors, we change the order from m to $m+1$, we add m to the set of marker orders, and we recalculate the flag counter.

Let us see how such transformation changes the flag counter on the path to the selected leaf. We will prove (by induction) that the previous value d and the new value d' of the flag counter in every node on this path satisfy $d' \geq \log_2 d$. In the selected leaf itself, the flag counter (being either 0 or 1) remains unchanged; we have $d' = d \geq \log_2 d$. Next, consider any proper ancestor of the selected node. Let k be the number of those of its children in which the flag counter was positive, plus the number of order- m flags placed in the considered node itself. Let also d_{\max} and d'_{\max} be the previous value and the new value of the flag counter in this child that is in the direction of the selected leaf. By construction, the flag counter in this child was maximal, which implies $k \cdot d_{\max} \geq d$, while by the induction assumption $d'_{\max} \geq \log_2 d_{\max}$. To d' we take the flag counter only from the special child, while for other children with positive flag counter we add 1, i.e., $d' = k - 1 + d'_{\max}$. Altogether we obtain $d' = k - 1 + d'_{\max} \geq k - 1 + \log_2 d_{\max} \geq \log_2 (k \cdot d_{\max}) \geq \log_2 d$, as required. \square

Proof of Lemma 5. We consider the base case when $P = (\lambda x.R)S$ and $Q = R[S/x]$; the general situation (redex being deeper in P) is easily reduced to this one. In the derivation of $\Gamma \vdash Q : \hat{\tau} \triangleright c$ we identify the set I of places (nodes) where we derive a type for S substituted for x . For $i \in I$, let $\Sigma_i \vdash S : \hat{\sigma}_i \triangleright d_i$ be the type judgment in i . We change the nodes in I into leaves, where we instead derive $\varepsilon[x \mapsto \{\hat{\sigma}_i\}] \vdash x : \hat{\sigma}_i \triangleright 0$. It should be clear that we can repair the rest of the derivation, by changing type environments, replacing S by x in λ -terms, and decreasing flag counters. In this way we obtain derivations of $\Sigma_i \vdash S : \hat{\sigma}_i \triangleright d_i$ for every $i \in I$, and a derivation of $\Sigma' \vdash R : \hat{\tau} \triangleright d$, where $\Sigma' = \Sigma[x \mapsto \{\hat{\sigma}_i \mid i \in I\}]$ with $\Sigma(x) = \emptyset$, and $\text{Split}(\Gamma \mid \Sigma, (\Sigma_i)_{i \in I})$, and $c = d + \sum_{i \in I} d_i$. To the latter type judgment we apply the (λ) rule, and then we merge it with the type judgments for S using the $(@)$ rule, which results in a derivation for $\Gamma \vdash P : \hat{\tau} \triangleright c$. We remark that different $i \in I$ may give identical type judgments for S (as long as the set of markers in $\hat{\sigma}_i$ is empty); this is not a problem. The $(@)$ rule requires that $\text{ord}(\hat{\sigma}_i) = \text{ord}(\lambda x.R)$; we have that $\text{ord}(\hat{\sigma}_i) = \text{ord}(\hat{\tau})$, and $\text{ord}(\hat{\tau}) = m = \text{ord}(\lambda x.R)$ by assumption. \square

5 Soundness

In this section we sketch the proof of the right-to-left implication of Theorem 2. We, basically, need to reverse the proof from the previous section. The following new fact is now needed.

Lemma 6. *If we can derive $\Gamma \vdash P : (m, F, M, \tau) \triangleright c$ with $m - 1 \notin M$ and $\text{ord}(P) \leq m - 1$, then $c = 0$.*

A simple inductive proof is based on the following idea: flags of order m are created only when a marker of order $m - 1$ is visible; the derivation itself (together with free variables) does not provide it ($m - 1 \notin M$), and the arguments, i.e. sets T_1, \dots, T_k in $\tau = T_1 \rightarrow \dots \rightarrow T_k \rightarrow o$, may provide only markers of order at most $\text{ord}(P) - 1 \leq m - 2$ (see the definition of a type), thus no flags of order m can be created.

We say that a λ -term of the form PQ is an application of order n when $\text{ord}(P) = n$, and that an $(@)$ rule is of order n if it derives a type for an application of order n . We can successively remove applications of the maximal order from a type derivation.

Lemma 7. *Suppose that $\varepsilon \vdash P : \hat{\rho}_m \triangleright c$ for $m > 0$ is derived by a derivation D in which the $(@)$ rule of order m is used n times. Then there exists Q such that $P \rightarrow_\beta Q$ and $\varepsilon \vdash Q : \hat{\rho}_m \triangleright c$ can be derived by a derivation D' in which the $(@)$ rule of order m is used less than n times.*

Recall from the definition of the type system that the $(@)$ rule of orders higher than m cannot be used while deriving a full type of order m . Thus in D we have type judgments only for subterms of P of order at most m (although P may also have subterms of higher orders), and in type environments we only have variables of order at most $m - 1$. In order to prove Lemma 7 we choose in P a subterm RS with $\text{ord}(R) = m$ such that there is a type judgment for RS in some nodes of D (at least one), but no descendants of those nodes use the $(@)$ rule of order m . Since R is of order m , it cannot be an application (then we would choose it instead of RS) nor a variable; thus $R = \lambda x.R'$. We obtain Q by reducing the redex $(\lambda x.R')S$; the derivation D' is obtained by performing a surgery on D similar to that in the proof of Lemma 5 (but in the opposite direction). Notice that every full type (m, F, M, τ) (derived for S) with nonempty M is used for exactly one appearance of x in the derivation for R' ; full types with empty M may be used many times, or not used at all, but thanks to Lemma 6 duplicating or removing the corresponding derivations for S does not change the flag counter. In the derivations for $R'[S/x]$ no $(@)$ rule of order m may appear, and the application RS disappears, so the total number of $(@)$ rules of order m decreases.

When all $(@)$ rules of order m are eliminated, we can decrease m .

Lemma 8. *Suppose that $\varepsilon \vdash P : \hat{\rho}_m \triangleright c$ for $m > 0$ is derived by a derivation D in which the $(@)$ rule of order m is not used. Then we can also derive $\varepsilon \vdash P : \hat{\rho}_{m-1} \triangleright c'$ for some $c' \geq c$.*

The proof is easy; we simply decrease the order m of all derived full types by 1, and we ignore flags of order m and markers of order $m - 1$. To obtain the inequality $c' \geq c$ we observe that when no $(@)$ rule of order m is used, the information about flags of order $m - 1$ goes only from descendants to ancestors, and thus every flag of order m is created because of a different flag of order $m - 1$.

By repeatedly applying the two above lemmata, out of a derivation of $\varepsilon \vdash P : \hat{\rho}_m \triangleright c$ we obtain a derivation of $\varepsilon \vdash Q : \hat{\rho}_0 \triangleright c'$, where $P \rightarrow_\beta^* Q$ and $c' \geq c$. Since $\hat{\rho}_0$ is of order 0, using the latter derivation it is easy to find in the already expanded part of Q (and thus in $\mathcal{L}(Q) = \mathcal{L}(P)$) a tree t such that $|t| = c' \geq c$.

6 Effectiveness

Finally, we show how Theorem 1 follows from Theorem 2, i.e., how given a λY -term P of complexity m we can check whether $\varepsilon \vdash P : \hat{\rho}_m \triangleright c$ can be derived for arbitrarily large c . We say that two type judgments are equivalent if they differ only in the value of the flag counter. Let us consider a set \mathcal{D} of all derivations

of $\varepsilon \vdash P : \hat{\rho}_m \triangleright c$ in which on each branch (i.e., each root-leaf path) there are at most three type judgments from every equivalence class, and among premisses of each (@) rule there is at most one type judgment from every equivalence class. These derivations use only type judgments $\Gamma \vdash Q : \hat{\tau} \triangleright d$ with Q being a subterm of P and with $\Gamma(x) \neq \emptyset$ only for variables x appearing in P . Since a finite λY -term, even when seen as an infinitary λ -term, has only finitely many subterms, this introduces a common bound on the height of all derivations in \mathcal{D} , and on their degree (i.e., on the maximal number of premisses of a rule). It follows that there are only finitely many derivations in \mathcal{D} , and thus we can compute all of them.

We claim that $\varepsilon \vdash P : \hat{\rho}_m \triangleright c$ can be derived for arbitrarily large c if and only if in \mathcal{D} there is a derivation in which on some branch there are two equivalent type judgments with different values of the flag counter (and the latter condition can be easily checked). Indeed, having such a derivation, we can repeat its fragment between the two equivalent type judgments, obtaining derivations of $\varepsilon \vdash P : \hat{\rho}_m \triangleright c$ with arbitrarily large c . We use here an additivity property of our type system: if out of $\Gamma \vdash Q : \hat{\tau} \triangleright d$ we can derive $\Gamma' \vdash Q' : \hat{\tau}' \triangleright d'$, then out of $\Gamma \vdash Q : \hat{\tau} \triangleright d + k$ we can derive $\Gamma' \vdash Q' : \hat{\tau}' \triangleright d' + k$, for every $k \geq -d$. Conversely, take a derivation of $\varepsilon \vdash P : \hat{\rho}_m \triangleright c$ for some large enough c . Suppose that some of its (@) rules uses two equivalent premisses. These premisses concern the argument subterm, which is of smaller order than the operator subterm, and thus of order at most $m - 1$. The set of marker orders in these premisses has to be empty, as the sets of marker orders from all premisses have to be disjoint. Thus, by Lemma 6, the flag counter in our two premisses is 0. In consequence, we can remove one of the premisses, without changing anything in the remaining part of the derivation, even the flag counters. In this way we clean the whole derivation, so that at the end among premisses of each (@) rule there is at most one type judgment from every equivalence class. The degree is now bounded, and at each node the flag counter grows only by a constant above the sum of flag counters from the children. Thus, if c is large enough, we can find on some branch two equivalent type judgments with different values of the flag counter. Then, for some pairs of equivalent type judgments, we remove the part of the derivation between these type judgments (and we adopt appropriately the flag counters in the remaining part). It is not difficult to perform this cleaning so that the resulting derivation will be in \mathcal{D} , and simultaneously on some branch there will remain two equivalent type judgments with different values of the flag counter.

7 Conclusions

In this paper, we have shown an approach for expressing quantitative properties of Böhm trees using an intersection type system, on the example of the finiteness problem. It is an ongoing work to apply this approach to the diagonal problem, which should give a better complexity than that of the algorithm from [6]. Another ongoing work is to obtain an algorithm for model checking Böhm trees with respect to the Weak MSO+U logic [3]. This logic extends Weak MSO by a new quantifier U, expressing that a subformula holds for arbitrarily large finite sets. Furthermore, it seems feasible that our methods may help in proving a pumping lemma for nondeterministic HORSEs.

References

- [1] Achim Blumensath (2008): *On the Structure of Graphs in the Caucal Hierarchy*. *Theor. Comput. Sci.* 400(1-3), pp. 19–45, doi:10.1016/j.tcs.2008.01.053.
- [2] Achim Blumensath (2013): *Erratum to "On the Structure of Graphs in the Caucal Hierarchy" [Theoret. Comput. Sci. 400 (2008) 19-45]*. *Theor. Comput. Sci.* 475, pp. 126–127, doi:10.1016/j.tcs.2012.12.044.

- [3] Mikołaj Bojańczyk & Szymon Toruńczyk (2012): *Weak MSO+U over Infinite Trees*. In Christoph Dürr & Thomas Wilke, editors: *29th International Symposium on Theoretical Aspects of Computer Science, STACS 2012, February 29th - March 3rd, 2012, Paris, France, LIPIcs 14*, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 648–660, doi:10.4230/LIPIcs.STACS.2012.648.
- [4] Christopher H. Broadbent, Arnaud Carayol, Matthew Hague & Olivier Serre (2012): *A Saturation Method for Collapsible Pushdown Systems*. In Artur Czumaj, Kurt Mehlhorn, Andrew M. Pitts & Roger Wattenhofer, editors: *Automata, Languages, and Programming - 39th International Colloquium, ICALP 2012, Warwick, UK, July 9-13, 2012, Proceedings, Part II, Lecture Notes in Computer Science 7392*, Springer, pp. 165–176, doi:10.1007/978-3-642-31585-5_18.
- [5] Christopher H. Broadbent & Naoki Kobayashi (2013): *Saturation-Based Model Checking of Higher-Order Recursion Schemes*. In Simona Ronchi Della Rocca, editor: *Computer Science Logic 2013 (CSL 2013), CSL 2013, September 2-5, 2013, Torino, Italy, LIPIcs 23*, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 129–148, doi:10.4230/LIPIcs.CSL.2013.129.
- [6] Lorenzo Clemente, Paweł Parys, Sylvain Salvati & Igor Walukiewicz (2016): *The Diagonal Problem for Higher-Order Recursion Schemes is Decidable*. In Martin Grohe, Eric Koskinen & Natarajan Shankar, editors: *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '16, New York, NY, USA, July 5-8, 2016, ACM*, pp. 96–105, doi:10.1145/2933575.2934527.
- [7] Lorenzo Clemente, Paweł Parys, Sylvain Salvati & Igor Walukiewicz (2015): *Ordered Tree-Pushdown Systems*. In Prahladh Harsha & G. Ramalingam, editors: *35th IARCS Annual Conference on Foundation of Software Technology and Theoretical Computer Science, FSTTCS 2015, December 16-18, 2015, Bangalore, India, LIPIcs 45*, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 163–177, doi:10.4230/LIPIcs.FSTTCS.2015.163.
- [8] Werner Damm (1982): *The IO- and OI-Hierarchies*. *Theor. Comput. Sci.* 20, pp. 95–207, doi:10.1016/0304-3975(82)90009-3.
- [9] Matthew Hague, Jonathan Kochems & C.-H. Luke Ong (2016): *Unboundedness and Downward Closures of Higher-Order Pushdown Automata*. In Rastislav Bodík & Rupak Majumdar, editors: *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016, ACM*, pp. 151–163, doi:10.1145/2837614.2837627.
- [10] Matthew Hague, Andrzej S. Murawski, C.-H. Luke Ong & Olivier Serre (2008): *Collapsible Pushdown Automata and Recursion Schemes*. In: *Proceedings of the Twenty-Third Annual IEEE Symposium on Logic in Computer Science, LICS 2008, 24-27 June 2008, Pittsburgh, PA, USA, IEEE Computer Society*, pp. 452–461, doi:10.1109/LICS.2008.34.
- [11] Alexander Kartzow & Paweł Parys (2012): *Strictness of the Collapsible Pushdown Hierarchy*. In Branislav Rován, Vladimiro Sassone & Peter Widmayer, editors: *Mathematical Foundations of Computer Science 2012 - 37th International Symposium, MFCS 2012, Bratislava, Slovakia, August 27-31, 2012. Proceedings, Lecture Notes in Computer Science 7464*, Springer, pp. 566–577, doi:10.1007/978-3-642-32589-2_50.
- [12] Teodor Knapik, Damian Niwiński & Paweł Urzyczyn (2002): *Higher-Order Pushdown Trees Are Easy*. In Mogens Nielsen & Uffe Engberg, editors: *Foundations of Software Science and Computation Structures, 5th International Conference, FOSSACS 2002. Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2002 Grenoble, France, April 8-12, 2002, Proceedings, Lecture Notes in Computer Science 2303*, Springer, pp. 205–222, doi:10.1007/3-540-45931-6_15.
- [13] Naoki Kobayashi (2009): *Types and Higher-Order Recursion Schemes for Verification of Higher-Order Programs*. In Zhong Shao & Benjamin C. Pierce, editors: *Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2009, Savannah, GA, USA, January 21-23, 2009, ACM*, pp. 416–428, doi:10.1145/1480881.1480933.
- [14] Naoki Kobayashi (2013): *Pumping by Typing*. In: *28th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2013, New Orleans, LA, USA, June 25-28, 2013, IEEE Computer Society*, pp. 398–407, doi:10.1109/LICS.2013.46.

- [15] Naoki Kobayashi, Kazuhiro Inaba & Takeshi Tsukada (2014): *Unsafe Order-2 Tree Languages Are Context-Sensitive*. In Anca Muscholl, editor: *Foundations of Software Science and Computation Structures - 17th International Conference, FOSSACS 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014, Proceedings, Lecture Notes in Computer Science* 8412, Springer, pp. 149–163, doi:10.1007/978-3-642-54830-7_10.
- [16] Naoki Kobayashi & C.-H. Luke Ong (2009): *A Type System Equivalent to the Modal Mu-Calculus Model Checking of Higher-Order Recursion Schemes*. In: *Proceedings of the 24th Annual IEEE Symposium on Logic in Computer Science, LICS 2009, 11-14 August 2009, Los Angeles, CA, USA*, IEEE Computer Society, pp. 179–188, doi:10.1109/LICS.2009.29.
- [17] C.-H. Luke Ong (2006): *On Model-Checking Trees Generated by Higher-Order Recursion Schemes*. In: *21th IEEE Symposium on Logic in Computer Science (LICS 2006), 12-15 August 2006, Seattle, WA, USA, Proceedings*, IEEE Computer Society, pp. 81–90, doi:10.1109/LICS.2006.38.
- [18] Paweł Parys (2016): *A Characterization of Lambda-Terms Transforming Numerals*. *J. Funct. Program.* 26, p. e12, doi:10.1017/S0956796816000113.
- [19] Paweł Parys (2012): *On the Significance of the Collapse Operation*. In: *Proceedings of the 27th Annual IEEE Symposium on Logic in Computer Science, LICS 2012, Dubrovnik, Croatia, June 25-28, 2012*, IEEE Computer Society, pp. 521–530, doi:10.1109/LICS.2012.62.
- [20] Paweł Parys (2017): *Intersection Types and Counting*. CoRR abs/1701.05303v1. Available at <http://arxiv.org/abs/1701.05303v1>.
- [21] Steven J. Ramsay, Robin P. Neatherway & C.-H. Luke Ong (2014): *A Type-Directed Abstraction Refinement Approach to Higher-Order Model Checking*. In Suresh Jagannathan & Peter Sewell, editors: *The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '14, San Diego, CA, USA, January 20-21, 2014*, ACM, pp. 61–72, doi:10.1145/2535838.2535873.
- [22] Sylvain Salvati & Igor Walukiewicz (2016): *Simply Typed Fixpoint Calculus and Collapsible Pushdown Automata*. *Mathematical Structures in Computer Science* 26(7), pp. 1304–1350, doi:10.1017/S0960129514000590.